Veritas Storage Foundation™ Cluster File System Release Notes

HP-UX

5.0 Maintenance Pack 2

symantec™

Veritas Storage Foundation™ Cluster File System Release Notes

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product Version: 5.0 MP2

Document version: 5.0MP2.0

Legal Notice

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's maintenance offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

For information about Symantec's Maintenance Programs, you can visit our Web site at the following URL:

http://www.symantec.com/techsupp/

Contacting Technical Support

Customers with a current maintenance agreement may access Technical Support information at the following URL:

http://www.symantec.com/business/support/index.jsp

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information
- Available memory, disk space, and NIC information
- Operating system

- Version and patch level

- Network topology

- Router, gateway, and IP address information

- Problem description:

  - Error messages and log files

  - Troubleshooting that was performed before contacting Symantec

  - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

http://www.symantec.com/techsupp/

Customer service

Customer service information is available at the following URL:

http://www.symantec.com/techsupp/

Customer Service is available to assist with the following types of issues:

- Questions regarding product licensing or serialization

- Product registration updates, such as address or name changes

- General product information (features, language availability, local dealers)

- Latest information about product updates and upgrades

- Information about upgrade assurance and maintenance contracts

- Information about the Symantec Buying Programs

- Advice about Symantec's technical support options

- Nontechnical presales questions

- Issues that are related to CD-ROMs or manuals

Maintenance agreement resources

If you want to contact Symantec regarding an existing maintenance agreement, please contact the maintenance agreement administration team for your region as follows:

Asia-Pacific and Japan	contractsadmin@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Additional enterprise services

Symantec offers a comprehensive set of services that allow you to maximize your investment in Symantec products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively.

Enterprise services that are available include the following:

Symantec Early Warning Solutions	These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.
Managed Security Services	These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Educational Services	Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about Enterprise services, please visit our Web site at the following URL:

http://www.symantec.com

Select your country or language from the site index.

Storage Foundation Cluster File System Release Notes

This document includes the following topics:

Introduction

This document provides information specific to the Maintenance Pack 2 (MP2) for Storage Foundation Cluster File System (SFCFS) 5.0.

■ For information on topics not covered in this document, such as SFCFS-specific software fixes and limitations, see the *Veritas Storage Foundation Cluster File System 5.0 Release Notes*.

■ For information on Storage Foundation, including software fixes and limitations for Veritas File System, Veritas Volume Replicator, and Veritas Volume Manager, see the *Veritas Storage Foundation 5.0 MP2 Release Notes*.

Read the following Veritas Technical Support TechNote for the latest information on updates, patches, and software issues for this release:

http://entsupport.symantec.com/docs/319349

System requirements

This section describes the system requirements for this release.

Software and hardware requirements

For information on hardware requirements, see the hardware compatibility list. The hardware compatibility list (HCL) is available at:

http://entsupport.symantec.com/docs/283161

For information on specific HA setup requirements, see the *Veritas Cluster Server Installation Guide*.

Supported HP-UX operating systems

This release of Veritas Storage Foundation, Veritas Storage Foundation Cluster File System, and Veritas Storage Foundation for Oracle can only be installed on a system running the September 2004 HP-UX 11i version 2.0 release or later on the PA-RISC or Itanium platforms. However, Symantec recommends to install the latest Fusion release (June 2008 or later) of HP-UX 11i version 2 operating system before upgrading to 5.0 Maintenance Pack 2.

To verify the operating system version

Use the swlist command as follows:

```
# swlist | grep HPUXBaseAux
HPUXBaseAux        B.11.23.0409 HP-UX Base OS Auxiliary
```

JFS must be installed on your system prior to installing any Veritas software.

To verify that JFS is installed

Use the `swlist` command as follows:

```
# swlist -l product JFS
JFS B.11.23      The Base VxFS File System
```

Required HP-UX patches

The 5.0 MP2 releases of Veritas Storage Foundation and Veritas Storage Foundation for Oracle require the following HP-UX patches.

> **Warning:** Install all the latest required HP-UX patches before you install Veritas Storage Foundation, Veritas Storage Foundation High Availability Solutions, Storage Foundation Cluster File System, or Veritas Volume Replicator. You can use the `swlist` command to determine whether the correct update and patches are installed. The installation procedure terminates if the correct patches are not found. Make sure that `EnableVXFS` bundle has revision B.11.23.04 or later after installing the latest patches.

HP may release patches that supersede the ones in this list. To verify that you have the latest HP-UX patches, go to the Symantec support website to view the relevant TechNote.

http://www.symantec.com/techsupp

Also, you can get the patches from Hewlett-Packard's Patch Database offered under the Maintenance and Support section of the HP Services & Support - IT Resource Center. HP's Patch Database provides fast, accurate searches for the latest recommended and superseded patches available for Veritas File System or Veritas Volume Manager.

Most of the above patches are available in the Feature11i bundle. The Feature11i bundle is available from HP software download site:

http://h20293.www2.hp.com/

Table 1-1   Required HP-UX patches

HP-UX Patch ID	Description
PHCO_32385	Enables `fscat`(1M).
PHCO_32387	Enables `getext`(1M).

Table 1-1   Required HP-UX patches *(continued)*

HP-UX Patch ID	Description
PHCO_32388	Enables setext(1M).
PHCO_32389	Enables vxdump(1M).
PHCO_32390	Enables vxrestore(1M).
PHCO_32391	Enables vxfsstat(1M).
PHCO_32392	Enables vxtunefs(1M).
PHCO_32393	Enables vxupgrade(1M).
PHCO_32488	Enables LIBC for VxFS 4.1 and later file systems.
PHCO_32523	Enhancement to quota(1) for supporting large uids.
PHCO_32524	Enhancement to edquota for supporting large uids.
PHCO_32551	Enhancement to quotaon/quotaoff for supporting large uids.
PHCO_32552	Enhancement to repquota for supporting large uids.
PHCO_32596	Enables df(1M).
PHCO_32608	Enables bdf(1M).
PHCO_32609	Enables fstyp(1M).
PHCO_32610	Enables mount(1M).
PHCO_32611	Fix fs_wrapper to accept "vxfs" from subtype.
PHCO_33238	swapon(1M) cumulative patch.
PHCO_34036	LVM commands patch.
PHCO_34208	SAM cumulative patch.
PHCO_34191	Cumulative libc patch.
PHSS_32228	LIBCL patch. If the patch is not installed, the VAILAgent process could terminate and LUN-to-disk mapping would fail.
PHSS_32231	FORTRAN I/O Library patch. If the patch is not installed, the VAILAgent process could terminate and LUN-to-disk mapping would fail.
PHSS_32674	Obam patch (backend for the SAM patch).

Table 1-1   Required HP-UX patches *(continued)*

HP-UX Patch ID	Description
PHKL_31500	Sept04 Base patch.
PHKL_32272	Changes to fix intermittent failures in `getacl/setacl`.
PHKL_32430	Changes to separate vxfs symbols from libdebug.a, so that symbols of VxFS 4.1 and later are easily available in q4/p4.
PHKL_32431	Changes to disallow mounting of a file system on a vnode having VNOMOUNT set. Enhancements for supporting quotas on large uids.
PHKL_33312	LVM Cumulative patch.
PHKL_34010	Cumulative VM patch.

In addition to the above patches the `EnableVXFS` bundle needs to be installed before installing the VxFS 5.0 MP2 file system. This bundle is an HP bundle and contains enhancements to various commands to understand the new disk layout Version 6 and later.

Table 1-2 lists the patches in the `EnableVXFS` bundle.

Table 1-2   EnableVXFS bundle

HP-UX Patch ID	Description
FsLibEnh	Enhancements to LIBC libraries to understand VxFS disk layout Version 6 and later.
DiskQuota-Enh	Enhancements to various quota related commands to support large uids.
FsCmdsEnh	Enhancements to the `mount` command to support VxFS 4.1.

For VEA functionality to work reliably, you must have HP-UX patches PHSS_33903 and PHSS_38526 (Itanium systems only) installed prior to installing Storage Foundation Cluster File System 5.0 Maintenance Pack 2.

Storage Foundation Cluster File System node requirements

All nodes in a Cluster File System must have the same operating system version and update level.

Component product release notes

In addition to reading these Storage Foundation Cluster File System Release Notes, review all component product release notes before installing the product.

The following component product release notes are included as PDF files on the software disc:

- *Veritas Storage Foundation Release Notes* (`sf_notes.pdf`)
- *Veritas Cluster Server Release Notes* (`vcs_notes.pdf`)

About the common product installer

The product installer is the recommended method to license and install the Veritas products. The installer also enables you to configure the product, verify preinstallation requirements, and view the product's description.

If you obtained a standalone Veritas product from an electronic download site, the single product download files do not contain the general product installer. Use the product installation script to install the product.

At most points during an installation, you can type b (back) to return to a previous section of the installation procedure. The back feature of the installation scripts is context-sensitive, so it returns to the beginning of a grouped section of questions. If an installation procedure hangs, use Control–c to stop and exit the program. After a short delay, the script exits.

Default responses are in parentheses. Press Return to accept the defaults.

Additional options are available for the common product installer.

Installing Storage Foundation Cluster File System using the common product installer

The product installer is the recommended method to license and install Storage Foundation Cluster File System.

For a simple initial installation on new system, you can use the following procedure.

The following sample procedure is based on the installation of a Veritas Storage Foundation Cluster File System HA cluster with two nodes: "system01" and "system02." If you are installing on standalone systems only, some steps are unnecessary, and these are indicated.

Default responses are enclosed by parentheses. Press Return to accept defaults.

To install the Storage Foundation Cluster File System

1   To install on multiple systems, set up the systems so that commands between systems execute without prompting for passwords or confirmations.

2   Load and mount the software disc.

3   Move to the top-level directory on the disc.

4   From this directory, type the following command to install if you are using the secure shell (ssh) utilities:

```
# ./installer
```

If you use the remote shell utilities to install on remote systems, additionally specify the -rsh option:

```
# ./installer -rsh
```

The sample installation assumes that ssh is used.

5   From the Installation menu, choose the I option for Install and enter the number for Veritas Storage Foundation Cluster File System. Press **Return**.

6   You are prompted to enter one or more system names to install SFCFS.

```
Enter the system names separted by spaces on which to install
SFCFS: system01 system02
```

7   During the initial system check, the installer verifies that communication between systems has been set up.

If the installer hangs or asks for a login password, stop the installer and set up ssh or rsh. Then run the installer again.

8   Enter the product license information.

Each system requires a product license before installation. License keys for additional product features should also be added at this time.

```
Enter a SFCFS license key for system01?
```

9   Enter y to accept another license key or enter n to proceed.

```
Do you want to enter another license key for system02?
[y,n,q] (n) n
```

10 You can choose to install required packages or all packages.

```
1)Required Veritas Storage Foundation Cluster File System depots -
1566 MB required
2)All Veritas Storage Foundation Cluster File System depots -
1623 MB required
3)Storage Foundation for Oracle RAC depots - 2029 MB required

Select the depots to be installed on all systems? [1-3,q,?] (3) 1
```

11 A list includes the items in the selected option. Press **Return** to continue.

12 Reboot all the nodes on which SFCFS is installed.

13 If you are installing SFCFS for the first time proceed to Configuring Storage Foundation Cluster File System section.

Configuring Storage Foundation Cluster File System

This section describes configuring Storage Foundation Cluster File System using the Veritas Storage Foundation Cluster File System product installer.

To configure the product, run the Veritas product installer or the appropriate installation script using the `-configure` option.

To configure Storage Foundation Cluster File System

1 To invoke the configuration, run the `installsfcfs` command with the `configure` option, as shown in this example:

```
# ./installsfcfs -configure
```

2 You are prompted to enter the system names (in the following example, "system01" and "system02") on which the software is to be installed. Enter the system name or names and then press Return.

```
Enter the system names separated by spaces on which to
configure SFCFS: system01 system02
```

3 During the initial system check, the installer checks that communication between systems has been set up.

The installer requires that ssh commands used between systems execute without prompting for passwords or confirmations. If the installer hangs or asks for a login password, stop the installer and run it again with the ssh configured for password free logins, or configure rsh and use the -rsh option.

4 The procedure checks system licensing, and you can enter additional licenses, if needed.

```
Checking system licensing

SFCFS license registered on system01

Do you want to enter another license key for system01? [y,n,q] (n) n
```

5 Any running SFCFS processes are stopped. Enter Return to continue.

6 Starting I/O Fencing in enabled mode requires manual intervention after SFCFS Configuration. I/O Fencing can be configured in disabled mode now and it does not require any manual intervention after SFCFS Configuration.

Determine at this time if you plan to configure I/O Fencing in enabled mode or disabled mode, as well as the number of network interconnects (NICS) required on your systems. If you configure I/O Fencing in enabled mode only a single NIC is required, though at least two is recommended.

Enter y or n for configuring I/O Fencing in enabled mode.

```
Will you be configuring I/O Fencing in enabled mode?
  [y,n,q,?] (y) n
```

See the *Storage Foundation Cluster File System Administrator's Guide* for more information.

7 No configuration changes are made to the systems until all configuration questions are completed and confirmed. Press Return to continue.

All systems are configured to create one cluster.

Enter the unique cluster name and Cluster ID number.

```
Enter the unique cluster name: [?]cluster2
Enter the unique Cluster ID number between 0-65535: [b,?] 76
```

8 The installer discovers the NICs available on the first system and reports them.

```
Discovering NICs on host1 ... discovered lan0 lan1 lan2 lan3 lan4 lan5
```

9 Enter private heartbeat NIC information for each host.

```
Enter the NIC for the first private heartbeat
link on host1: [b,?] lan2
Would you like to configure a second private heartbeat
link? [y,n,q,b,?] (y) y
Enter the NIC for the second private heartbeat
link on host1: [b,?] lan3

Would you like to configure a third private heartbeat
link? [y,n,q,b,?] (n) n
Do you want to configure an additional low priority
heartbeat link? [y,n,q,b,?] (n) n
Are you using the same NICs for private heartbeat links
on all systems? [y,n,q,b,?] (y) y
```

> **Warning:** When answering y, be sure that the same NICs are available on each system; the installer may not verify this.

Notice that in this example, lan0 is not selected for use as a private heartbeat NIC because it already in use as the public network interface.

10 A summary of the information you entered is given. When prompted, confirm that the information is correct.

```
Is this information correct? [y,n,q]
```

If the information is correct, enter y. If the information is not correct, enter n. The installer prompts you to enter the information again.

11 The enclosure-based naming scheme (rather than disk device naming) is a feature of Veritas Volume Manager. You can reference disks using a symbolic name that is more meaningful than the operating system's normal device access name.

See the *Veritas Volume Manager Administrator's Guide*

```
Do you want to set up the enclosure-based naming scheme?
   [y,n,q,?] (n) n
```

12 You are now given the option of specifying the default name of a disk group that is to be assumed by Veritas Volume Manager commands if a disk group is not otherwise specified.

Enter **n** if you do not want to specify the name of the default disk group at this time. You can set the name of the default disk group after installation by running the `vxdctl defaultdg` *diskgroup* command on a system.

See the `vxdctl (1M)` manual page and the *Veritas Volume Manager Administrator's Guide* for more information.

If you specify the name of a default disk group, this step does not create the disk group. After installation, you can use the `vxdiskadm` command to create the disk group.

```
Do you want to set up a default disk group for each system?
 [y,n,q,?] (y) y
```

13 If you responded `y`, then enter the information for the default disk group name.

```
Will you specify one disk group name for all eligible
systems? [y,n,q,?] (y) y

Specify a default disk group name for all systems. [?] diskgroup001
```

14 Validate the default disk group information, and press Return.

15 You may be prompted to verify the fully qualified hostname of the systems.

```
Is the fully qualified hostname of system
 "system01" = system01.veritas.com"? [y,n,q] (y)
```

16 Enter `y` to accept the fully qualified domain name.

```
Is the fully qualified hostname of system
 "system02" = system02.veritas.com"? [y,n,q]
```

17 The cluster systems can be configured as a Storage Foundation Manager managed host or a stand-alone host.

Several prerequisites are necessary to configure the system as a Storage Foundation Manager managed host.

```
Enable Storage Foundation Management Server Management?
    [y,n,q] (y) n
```

18 The Veritas Storage Foundation Cluster File System software is verified and configured.

Check the log file, if needed, to confirm the configuration.

```
Configuration log files, summary file, and response file are saved at:

        /opt/VRTS/install/logs/installer-****
```

Upgrading Storage Foundation Cluster File System and High Availability

This section contains procedures for the Veritas Storage Foundation Cluster File System and High Availability upgrade.

About upgrading Storage Foundation Cluster File System and High Availability

Perform the procedures in the following sections to upgrade Storage Foundation Cluster File System and High Availability. You can perform an upgrade to Storage Foundation using the Veritas product installer or product installation script if you already have Storage Foundation Cluster File System installed.

See "Upgrading Storage Foundation Cluster File System and High Availability software from a release prior to 5.0" on page 20.

See "Upgrading Storage Foundation Cluster File System and High Availability software from a release 5.0 or later" on page 41.

Upgrade paths for Veritas Storage Foundation Cluster File System and High Availability

Table 1-3 shows the upgrade paths for Veritas Storage Foundation Cluster File System.

Table 1-3   Upgrade paths

From	Upgrade to	Tasks
Storage Foundation Cluster File System 3.5 (formerly known as SANPoint Foundation Suite 3.5)	Storage Foundation Cluster File System 5.0 MP2	See "Upgrading from 3.5 to 5.0 MP2" on page 23.
Storage Foundation Cluster File System 3.5 Update 1 (formerly known as SANPoint Foundation Suite 3.5 Update 1)	Storage Foundation Cluster File System 5.0 MP2	See "Upgrading from 3.5 to 5.0 MP2" on page 23.
Storage Foundation Cluster File System 3.5 Update 2 (formerly known as SANPoint Foundation Suite 3.5 Update 2)	Storage Foundation Cluster File System 5.0 MP2	See "Upgrading from 3.5 to 5.0 MP2" on page 23.
Storage Foundation Cluster File System 3.5 Update 3 (formerly known as SANPoint Foundation Suite 3.5 Update 3)	Storage Foundation Cluster File System 5.0 MP2	See "Upgrading from 3.5 to 5.0 MP2" on page 23.
Storage Foundation Cluster File System 3.5 Update 4 (formerly known as SANPoint Foundation Suite 3.5 Update 4)	Storage Foundation Cluster File System 5.0 MP2	See "Upgrading from 3.5 to 5.0 MP2" on page 23.

Table 1-3   Upgrade paths *(continued)*

From	Upgrade to	Tasks
Storage Foundation Cluster File System 4.1	Storage Foundation Cluster File System 5.0 MP2	See "Upgrading from 4.1 to 5.0 MP2" on page 31.
Storage Foundation Cluster File System 4.1 MP1	Storage Foundation Cluster File System 5.0 MP2	See "Upgrading from 4.1 to 5.0 MP2" on page 31.
Storage Foundation Cluster File System 4.1 MP2	Storage Foundation Cluster File System 5.0 MP2	See "Upgrading from 4.1 to 5.0 MP2" on page 31.
Storage Foundation Cluster File System 5.0 on HP-UX 11i version 2	Storage Foundation Cluster File System 5.0 MP2	See "Upgrading Storage Foundation Cluster File System and High Availability software from a release 5.0 or later" on page 41.
Storage Foundation Cluster File System 5.0 MP1 on HP-UX 11i version 2	Storage Foundation Cluster File System 5.0 MP2	See "Upgrading Storage Foundation Cluster File System and High Availability software from a release 5.0 or later" on page 41.

Upgrading Storage Foundation Cluster File System and High Availability software from a release prior to 5.0

This section contains procedures for the Veritas Storage Foundation Cluster File System and High Availability upgrade.

Planning the upgrade

Complete the following tasks in advance of upgrading:

- Review the *Veritas Storage Foundation Release Notes* for any late-breaking information on upgrading your system.

- Be sure that the administrator doing the upgrade has root access and a working knowledge of system administration.

- Schedule sufficient outage time for the upgrade.

- Make sure that the prerequisite patches required for SFCFS 5.0 MP2 are accessible.

Preparing the system and backing up files before upgrading

Before upgrading an installed Veritas Storage Foundation Cluster File System, preserve the existing configuration information.

To preserve the existing configuration information, perform the following actions:

- Make a record of the mount points for VxFS file systems and VxVM volumes that are defined in the `/etc/fstab` file. You will need to recreate these entries in the `/etc/fstab` file on the freshly installed system.

- Before upgrading, ensure that you have made backups of all data that you want to preserve. In particular, you will need the information in files such as `/etc/fstab`. You should also run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands, and record the output from these. You may need this information to reconfigure your system after the upgrade.

- Use the `vxlicrep` command to make a record of the currently installed Veritas licenses.

- Back up the configuration files.

  ```
  # cp -r /etc/VRTSvcs/conf/ backupdirectory/
  ```

- If Veritas Volume Replicator (VVR) is configured, do the following steps in the order shown:

  - Make sure that the disk groups that contain RVGs are at least at disk group version 110.

    ```
    # vxdg list diskgroup
    ```

  - Make sure the size of the SRL volume is greater than 110 MB.
    Refer to the *Veritas Volume Replicator Administrator's Guide*.

  - Stop all the applications involved in replication. For example, if a data volume contains a file system, unmount it.

  - Verify that all of the Primary RLINKs are up to date:

    ```
    # vxrlink -g diskgroup status rlink_name
    ```

Note: Do not continue until the Primary RLINKs are up to date.

Overview of procedures

There are two ways to upgrade cluster nodes to the latest version of Storage Foundation Cluster File System: phased and full.

Note: If VVR is configured, phased upgrade is not supported. We recommend that the secondary cluster be upgraded before the primary cluster in the RDS.

The upgrade procedures apply to both the phased and full upgrade procedures unless otherwise noted. Occasionally, steps differ between the two procedures. Screen output is also common between both procedures unless otherwise noted.

Note: Both procedures automatically uninstall the previous version of the software.

Phased upgrade

A phased upgrade minimizes downtime by upgrading portions of the cluster, one at a time. Although the entire cluster is offline for a shorter period than a full upgrade, this method requires command-line interaction and some manual configuration. Each phase of the phased upgrade should be performed on more than one node of the cluster. Cluster will be offline only during the last two bullets below for a phased upgrade.

The stages of the phased upgrade procedure are:

- Select two or more nodes to upgrade.

- Install the new version.

- Shut down VCS on remaining non-upgraded nodes.

- Modify the configuration information in the main.cf file.

- Install the new version on each remaining node and reboot them.

Note: A phased upgrade should not be performed from one of the nodes in the cluster.

Full upgrade

A full upgrade upgrades the product on the entire cluster and the cluster remains offline for the duration of the procedure. Minimal command-line interaction and some manual configuration are required.

The stages of the full upgrade procedure are:

■ Install the new version on all the nodes.

■ Modify the configuration information in the `main.cf` file.

■ Bring up the cluster.

Upgrading from 3.5 to 5.0 MP2

SFCFS can be upgraded from 3.5 to 5.0 MP2 using phased or full upgrade procedure.

Phased upgrade

Following procedure assumes a 4 node cluster system01, system02, system03, system04 where system01 and system02 are initially upgraded and rest of the cluster is brought up later.

To perform a phased upgrade

1 Log in as superuser.

2 Select one or more nodes to upgrade, say system01 and system02.

3 Insert the appropriate software disc into your system's DVD drive.

4 Determine the block device file for the DVD drive:

```
# ioscan -fnC disk
```

   Make a note of the device file as it applies to your system.

5 Create a directory in which to mount the software disc and mount the disc using the appropriate drive name. For example:

```
# mkdir -p /dvdrom
# /usr/sbin/mount -F cdfs /dev/dsk/c3t2d0 /dvdrom
```

6 Change to the top-level directory on the disc:

```
# cd /dvdrom
```

7 Offline all SFCFS resources on nodes selected in step 2 by running the following commands on one of the cluster nodes.

```
# hagrp -offline service_group -sys system01
# hagrp -offline service_group -sys system02
```

where *service_group* is the service group displayed by the `hagrp -dep cvm` command.

8 Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctrl status
```

9 If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctrl stop
```

10 Remove the VRTScavf and VRTSglm 3.5 packages from these nodes and run the following commands on system01 and system02.

```
# hastop -local
# swremove VRTScavf VRTSglm
```

11 Uninstall VCS 3.5 from system01 and system02. Run the following commands from one of the nodes.

See the *Veritas Cluster Server Installation Guide*.

```
# cd /opt/VRTSvcs/install
# ./uninstallvcs
```

Ignore any errors from the `uninstallvcs` script and proceed with the uninstallation of VCS.

12 Reboot system01 and system02 by running the following command on both nodes:

```
# /usr/sbin/shutdown -r now
```

13 Start VCS on the nodes in the cluster on which you did not uninstall VCS:

```
# /opt/VRTS/bin/hastart
```

14 Unload the `llt` module:

```
# kmadmin -U llt
```

15 Uninstall the `llt` and `gab` packages and patches from the nodes:

```
# swremove -x autoreboot=true VRTSllt VRTSgab PVKL_03675 \
PVKL_03674
```

16 Reboot the nodes.

17 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.

18 If system partitions other than `/stand` have disk layout Version 2 or 3, upgrade those partitions to disk layout Version 5 using the `vxupgrade`command. Partitions with disk layout Version 4 need not be upgraded.

19 Upgrade the operating system from HP-UX 11i Version 1 to HP-UX 11i Version 2 on system01 and system02.

See the HP-UX Operating System documentation.

20 Install all the prerequisite patches on system01 and system02.

See "Required HP-UX patches" on page 9.

21 Install SFCFS 5.0 MP2 on system01 and system02 and reboot these nodes.

See "Installing Storage Foundation Cluster File System using the common product installer" on page 12.

Do not configure SFCFS after reboot.

22 Check if there are frozen CVM and SFCFS groups from one of the non-upgraded nodes, enter:

```
# /opt/VRTSvcs/bin/hastatus -sum
```

If the groups are frozen, unfreeze CVM and SFCFS groups using the following commands for each group:

```
# haconf -makerw
# hagrp -unfreeze cvm -persistent
# hagrp -unfreeze service_group -persistent
# haconf -dump -makero
```

where *service_group* is the service group displayed by the `hagrp -dep cvm` command.

Run the following command on all the non-upgraded nodes, enter:

```
# hastop -local
```

23 Start vxfen on system01 and system02. vxfen can be started either in disable or enable mode. For starting vxfen in disabled mode, run the following commands

```
# echo vxfen_mode=disabled > /etc/vxfenmode
# /sbin/init.d/vxfen start
```

See the *Veritas Cluster Server Installation Guide*.

24 Change the configuration files by running the following commands on one of the upgraded nodes, say system01.

```
# /opt/VRTS/bin/hastart
# /opt/VRTS/bin/haconf -makerw
# hagrp -unfreeze cvm -persistent
# hagrp -unfreeze service_group -persistent
# /opt/VRTS/bin/hatype -add CVMVxconfigd
# /opt/VRTS/bin/hatype -modify CVMVxconfigd Operations \  OnOnly
# /opt/VRTS/bin/hares -add cvm_vxconfigd CVMVxconfigd cvm
# /opt/VRTS/bin/hares -modify cvm_vxconfigd Enabled 1
# /opt/VRTS/bin/hares -delete qlogckd
# /opt/VRTS/bin/haconf -dump -makero
# /opt/VRTS/bin/hastop -all -force
```

If you have configured the VCS Cluster Manager (Web Console), complete the following to modify the /etc/VRTSvcs/conf/config/main.cf file.

- Remove VRTSweb:

```
Process VRTSweb (
        PathName = "/opt/VRTSvcs/bin/haweb"
        Arguments = "10.129.96.64 8181"
        )
```

- Replace it with:

```
 VRTSWebApp VCSweb (
        Critical =0
        AppName = vcs
        InstallDir = "/opt/VRTSweb/VERITAS"
        TimeForOnline = 5
        )
```

■ Add the NIC resource in the ClusterService group. For example, where the name of the NIC resource is named csgnic and the public NIC device is hme0, add:

```
NIC csgnic (
     Device = hme0
```

■ Add new dependencies for the new resources in the ClusterService group. For example, using the names of the VRTSWebApp, NotifierMngr, IP, and NIC resources, enter lines that resemble:

```
VCSweb requires webip
  ntfr requires csgnic
  webip requires csgnic
```

25 Verify the syntax of the `/etc/VRTSvcs/conf/config/main.cf` file by running the following commands on system01:

```
# cd /etc/VRTSvcs/conf/config
# /opt/VRTS/bin/hacf -verify .
```

26 Start VCS on all the upgraded nodes. Run the following command on system01 and system02.

```
# /opt/VRTS/bin/hastart
# echo 3.5 > /opt/VRTS/install/.SFCFS.upgrade
```

27 Configure SFCFS on system01 and system02.

See "Configuring Storage Foundation Cluster File System" on page 14.

VCS configuration files are not changed during this configuration.

28 Upgrade file systems to proper disk layout version.

See "Upgrading the disk layout versions" on page 40.

29 Verify that all the file systems are working properly and data is intact.

See `cfsmount`(1M).

30 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctrl status
```

31 If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctrl start
```

32 Run 10 through 21, 23, 26, 27 and 29 on rest of the nodes to be upgraded.

Full upgrade

Following procedure assumes a 4 node cluster system01, system02, system03, system04 where all nodes are simultaneously upgraded from 3.5 to 5.0 MP2.

To perform a full upgrade

1 Log in as superuser.

2 Insert the appropriate software disc into your system's DVD drive.

3 Determine the block device file for the DVD drive:

```
# ioscan -fnC disk
```

Make a note of the device file as it applies to your system.

4 Create a directory in which to mount the software disc and mount the disc using the appropriate drive name. For example:

```
# mkdir -p /dvdrom
# /usr/sbin/mount -F cdfs /dev/dsk/c3t2d0 /dvdrom
```

5 Change to the top-level directory on the disc:

```
# cd /dvdrom
```

6 Offline all SFCFS resources on all nodes by running the following commands on one of the cluster nodes.

```
# hagrp -offline service_group -sys system01
# hagrp -offline service_group -sys system02
# hagrp -offline service_group -sys system03
# hagrp -offline service_group -sys system04
```

where *service_group* is the service group displayed by the hagrp -dep cvm command.

7 Remove the VRTScavf and VRTSglm 3.5 packages from these nodes and run the following command on all the systems.

```
# hastop -local
# swremove VRTScavf VRTSglm
```

8 Uninstall VCS 3.5 from all the nodes. Run the following commands from one of the nodes.

See the *Veritas Cluster Server Installation Guide.*

```
# cd /opt/VRTSvcs/install
# ./uninstallvcs
```

Ignore any errors from the uninstallvcs script and proceed with the uninstall of VCS.

9 Reboot the nodes.

10 Unload the llt module:

```
# kmadmin -U llt
```

11 Uninstall the llt and gab packages and patches from the nodes:

```
# swremove -x autoreboot=true VRTSllt VRTSgab PVKL_03675 \
PVKL_03674
```

12 Reboot the nodes.

13 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.

14 If system partitions other than /stand have disk layout Version 2 or 3, upgrade those partitions to disk layout Version 5 using the vxupgradecommand. Partitions with disk layout Version 4 need not be upgraded.

15 Upgrade the operating system from HP-UX 11i Version 1 to HP-UX 11i Version 2 on all the nodes.

See the HP-UX Operating System documentation.

16 Install all the prerequisite patches on all the nodes.

See "Required HP-UX patches" on page 9.

17 Install SFCFS 5.0 MP2 and reboot all the nodes.

See "Installing Storage Foundation Cluster File System using the common product installer" on page 12.

Do not configure SFCFS after reboot.

18 Start `vxfen` on all the nodes. `vxfen` can be started either in disable or enable mode. For starting `vxfen` in disabled mode, run the following commands:

```
# echo vxfen_mode=disabled > /etc/vxfenmode
# /sbin/init.d/vxfen start
```

See the *Veritas Cluster Server Installation Guide*.

19 Change the configuration files by running the following commands from one of the nodes.

```
# /opt/VRTS/bin/hastart
# /opt/VRTS/bin/haconf -makerw
# /opt/VRTS/bin/hatype -add CVMVxconfigd
# /opt/VRTS/bin/hatype -modify CVMVxconfigd Operations \
    OnOnly
# /opt/VRTS/bin/hares -add cvm_vxconfigd CVMVxconfigd cvm
# /opt/VRTS/bin/hares -modify cvm_vxconfigd Enabled 1
# /opt/VRTS/bin/hares -delete qlogckd
# /opt/VRTS/bin/haconf -dump -makero
# /opt/VRTS/bin/hastop -all -force
```

If you have configured the VCS Cluster Manager (Web Console), complete the following to modify the `/etc/VRTSvcs/conf/config/main.cf` file.

■ Remove VRTSweb:

```
Process VRTSweb (
        PathName = "/opt/VRTSvcs/bin/haweb"
        Arguments = "10.129.96.64 8181"
         )
```

■ Replace it with:

```
VRTSWebApp VCSweb (
        Critical =0
        AppName = vcs
        InstallDir = "/opt/VRTSweb/VERITAS"
        TimeForOnline = 5
         )
```

■ Add the NIC resource in the ClusterService group. For example, where the name of the NIC resource is named csgnic and the public NIC device is hme0, add:

```
NIC csgnic (
        Device = hme0
```

■ Add new dependencies for the new resources in the ClusterService group. For example, using the names of the VRTSWebApp, NotifierMngr, IP, and NIC resources, enter lines that resemble:

```
VCSweb requires webip
ntfr requires csgnic
webip requires csgnic
```

20 Verify the syntax of the /etc/VRTSvcs/conf/config/main.cf file by running the following command on system01:

```
# cd /etc/VRTSvcs/conf/config
# /opt/VRTS/bin/hacf -verify .
```

21 Run the following command on all the nodes to start VCS.

```
# /opt/VRTS/bin/hastart
# echo 3.5 > /opt/VRTS/install/.SFCFS.upgrade
```

22 Configure SFCFS on system01 and system02.

See "Configuring Storage Foundation Cluster File System" on page 14.

VCS configuration files are not changed during this configuration.

23 Upgrade file systems to proper disk layout version.

See "Upgrading the disk layout versions" on page 40.

24 Verify that all the file systems are working properly and data is intact.

See cfsmount(1M) manual page.

Upgrading from 4.1 to 5.0 MP2

SFCFS can be upgraded from 4.1 to 5.0 MP2 using phased or full upgrade procedure.

Phased upgrade

Following procedure assumes a 4 node cluster system01, system02, system03, system04 where system01 and system02 are initially upgraded and rest of the cluster is brought up later.

To perform a phased upgrade

1   Log in as superuser.

2   Select one or more nodes to upgrade, say system01 and system02.

3   Insert the appropriate software disc into your system's DVD drive.

4   Determine the block device file for the DVD drive:

    # **ioscan -fnC disk**

    Make a note of the device file as it applies to your system.

5   Create a directory in which to mount the software disc and mount the disc
    using the appropriate drive name. For example:

    # **mkdir -p /dvdrom**
    # **/usr/sbin/mount -F cdfs /dev/dsk/c3t2d0 /dvdrom**

6   Change to the top-level directory on the disc:

    # **cd /dvdrom**

7   Install all the prerequisite patches on system01 and system02.

    See "Required HP-UX patches" on page 9.

    If this step is being performed in any phase other than the first phase of the
    upgrade, the llthosts need to be replace to prevent vxfen from starting after
    reboot by running the following commands.

    ■ Move /etc/llthosts to /etc/llthosts.bak on all the nodes to be
      upgraded.

      # **mv /etc/llthosts /etc/llthosts.bak**

    ■ Install all the prerequisite patches and reboot the machines.

    ■ Move /etc/llthosts to /etc/llthosts.bak on all the nodes to be
      upgraded.

      # **mv /etc/llthosts.bak /etc/llthosts**

8 Offline all SFCFS resources on nodes selected in step 2 by running the following commands on one of the cluster nodes.

```
# hagrp -offline service_group -sys system01
# hagrp -offline service_group -sys system02
```

where *service_group* is the service group displayed by the `hagrp -dep cvm` command.

9 Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctrl status
```

10 If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctrl stop
```

11 Install SFCFS 5.0 MP2 on system01 and system02 and reboot these nodes.

See "Installing Storage Foundation Cluster File System using the common product installer" on page 12.

Do not configure SFCFS after reboot.

12 Check if there are frozen CVM and SFCFS groups from one of the non-upgraded nodes, enter:

```
# /opt/VRTSvcs/bin/hastatus -sum
```

If the groups are frozen, unfreeze CVM and SFCFS groups using the following commands for each group:

```
# haconf -makerw
# hagrp -unfreeze cvm -persistent
# hagrp -unfreeze service_group -persistent
# haconf -dump -makero
```

where *service_group* is the service group displayed by the `hagrp -dep cvm`command.

Run the following command on all the non-upgraded nodes, enter:

```
# hastop -local
```

13 Start `vxfen` on system01 and system02. `vxfen` can be started either in disable or enable mode. For starting `vxfen` in disabled mode, run the following commands:

```
# echo vxfen_mode=disabled > /etc/vxfenmode
# /sbin/init.d/vxfen start
```

See the *Veritas Cluster Server Installation Guide*.

14 Change the configuration files by running the following commands on one of the upgraded nodes. For example, system01.

```
# /opt/VRTS/bin/hastart
# /opt/VRTS/bin/haconf -makerw
# hagrp -unfreeze cvm -persistent
# hagrp -unfreeze service_group -persistent
# /opt/VRTS/bin/hares -delete qlogckd
# /opt/VRTS/bin/haconf -dump -makero
# /opt/VRTS/bin/hastop -all -force
```

where *service_group* is the service group displayed by the `hagrp -dep cvm` command.

15 If you have configured the VCS Cluster Manager (Web Console), complete the following to modify the `/etc/VRTSvcs/conf/config/main.cf` file.

- Remove VRTSweb:

```
Process VRTSweb (
        PathName = "/opt/VRTSvcs/bin/haweb"
        Arguments = "10.129.96.64 8181"
        )
```

- Replace it with:

```
VRTSWebApp VCSweb (
        Critical =0
        AppName = vcs
        InstallDir = "/opt/VRTSweb/VERITAS"
        TimeForOnline = 5
        )
```

- Add the NIC resource in the ClusterService group. For example, where the name of the NIC resource is named csgnic and the public NIC device is hme0, add:

```
NIC csgnic (
      Device = hme0
```

- Add new dependencies for the new resources in the ClusterService group. For example, using the names of the VRTSWebApp, NotifierMngr, IP, and NIC resources, enter lines that resemble:

```
VCSweb requires webip
    ntfr requires csgnic
    webip requires csgnic
```

16 Verify the syntax of the /etc/VRTSvcs/conf/config/main.cf file by running the following commands on system01:

```
# cd /etc/VRTSvcs/conf/config
# /opt/VRTS/bin/hacf -verify .
```

17 Start VCS on all the upgraded nodes. Run the following command on system01 and system02.

```
# /opt/VRTS/bin/hastart
```

18 Configure SFCFS on system01 and system02.

See "Configuring Storage Foundation Cluster File System" on page 14.

VCS configuration files are not changed during this configuration.

19 Upgrade file systems to proper disk layout version.

See "Upgrading the disk layout versions" on page 40.

20 Verify that all the file systems are working properly and data is intact.

See cfsmount(1M) manual page.

21 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctrl status
```

22 If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctrl start
```

23 Run 7, 11, and 18 on rest of the nodes to be upgraded.

Full upgrade

Following procedure assumes a 4 node cluster system01, system02, system03, system04 where all nodes are simultaneously upgraded from 4.1 to 5.0 MP2.

To perform a full upgrade

1   Log in as superuser.

2   Insert the appropriate software disc into your system's DVD drive.

3   Determine the block device file for the DVD drive:

```
# ioscan -fnC disk
```

Make a note of the device file as it applies to your system.

4   Create a directory in which to mount the software disc and mount the disc using the appropriate drive name. For example:

```
# mkdir -p /dvdrom
# /usr/sbin/mount -F cdfs /dev/dsk/c3t2d0 /dvdrom
```

5   Change to the top-level directory on the disc:

```
# cd /dvdrom
```

6   Install all the prerequisite patches on all the nodes.

See "Required HP-UX patches" on page 9.

7   Offline all SFCFS resources on all nodes by running the following commands on one of the cluster nodes.

```
# hagrp -offline service_group -sys system01
# hagrp -offline service_group -sys system02
# hagrp -offline service_group -sys system03
# hagrp -offline service_group -sys system04
```

where *service_group* is the service group displayed by the hagrp -dep cvm command.

8   Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctrl status
```

9   If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctrl stop
```

10 Install SFCFS on all the nodes.

- From this directory, type the following command to install if you are using the secure shell (`ssh`) utilities:

  ```
  # ./installer
  ```

- If you use the remote shell utilities to install on remote systems, additionally specify the `-rsh` option:

  ```
  # ./installer -rsh
  ```

The sample installation assumes that `-rsh` is used.

Note: Do not configure SFCFS after reboot.

---

11 From the Installation menu, choose the `I` option for Install and enter the number for Veritas Storage Foundation Cluster File System. Press **Return**.

12 Enter `y` to upgrade to version 5.0 on these systems using the current configuration.

```
Do you want to upgrade to version 5.0 on these systems using
the current configuration? [y,n,q,?] (y)
```

13 Enter `y` to accept another license key or enter `n` to proceed.

```
Do you want to enter another license key for system02?
[y,n,q] (n)
```

14 Enter `1` to install SFCFS depots.

```
1)  Install all Veritas Storage Foundation Cluster File System
depots - 9340 MB required
2)  Install Storage Foundation for Oracle RAC depots - 9987 MB
required

Select the depots to be installed on all systems? [1-2,q] (2) 1
```

15 A list includes the items in the selected option. Press **Return** to continue.

16 Enter `y` to upgrade SFCFS.

```
Are you sure you want to upgrade SFCFS? [y,n,q] (y)
```

17 Enter `y` to begin the SFCFS upgrade.

```
Are you ready to begin the Veritas Storage Foundation Cluster
File System upgrade at this time [y,n,q] (y)
```

18 Reboot all the nodes.

```
# /usr/sbin/shutdown -r now
```

19 Start `vxfen` on all the nodes. `vxfen` can be started either in disable or enable mode. For starting vxfen in disabled mode, run the following commands:

```
# echo vxfen_mode=disabled > /etc/vxfenmode
# /sbin/init.d/vxfen start
```

See the *Veritas Cluster Server Installation Guide* for information regarding starting `vxfen` in enabled mode.

20 Change the configuration files by running the following commands from one of the nodes.

```
# /opt/VRTS/bin/hastart
# /opt/VRTS/bin/haconf -makerw
# /opt/VRTS/bin/hares -delete qlogckd
# /opt/VRTS/bin/haconf -dump -makero
# /opt/VRTS/bin/hastop -all -force
```

21 If you have configured the VCS Cluster Manager (Web Console), complete the following to modify the `/etc/VRTSvcs/conf/config/main.cf` file.

- Remove VRTSweb:

```
Process VRTSweb (
  PathName = "/opt/VRTSvcs/bin/haweb"
  Arguments = "10.129.96.64 8181"
  )
```

- Replace it with:

```
VRTSWebApp VCSweb (
            Critical =0
            AppName = vcs InstallDir = "/opt/VRTSweb/VERITAS"
            TimeForOnline = 5
            )
```

■ Add the NIC resource in the ClusterService group. For example, where the
name of the NIC resource is named csgnic and the public NIC device is
hme0, add:

```
NIC csgnic (
        Device = hme0
```

■ Add new dependencies for the new resources in the ClusterService group.
For example, using the names of the VRTSWebApp, NotifierMngr, IP, and
NIC resources, enter lines that resemble:

```
VCSweb requires webip
    ntfr requires csgnic
    webip requires csgnic
```

22 Verify the syntax of the /etc/VRTSvcs/conf/config/main.cf file by running
the following commands on system01:

```
# cd /etc/VRTSvcs/conf/config
# /opt/VRTS/bin/hacf -verify .
```

23 Run the following command on all the nodes to start VCS.

```
# /opt/VRTS/bin/hastart
```

24 Configure SFCFS on all the nodes.

Invoke the configuration, run the installsfcfs command with the
-configure option, as shown in this example:

```
# ./installsfcfs -configure -rsh
```

25 If you enter y then the current configuration will be used to configure SFCFS.
If you enter n then you can have a different configuration to configure SFCFS.

```
Do you want to configure SFCFS from these systems? [y,n,q] (y)
```

26 Enter y to accept another license key or enter n to proceed.

```
Do you want to enter another license key for system02?
```

27 Any running SFCFS processes are stopped. Press **Return** to continue.

28 The cluster systems can be configured as a Storage Foundation Manager managed host or a stand-alone host. Several prerequisites are necessary to configure the system as a Storage Foundation Manager managed host.

```
Enable Storage Foundation Management Server Management?
[y,n,q] (y) n
```

29 Enter y to start SFCFS processes.

```
Do you want to start Veritas Storage Foundation Cluster File System
processes now? [y,n,q] (y)
```

30 Upgrade file systems to proper disk layout version.

See "Upgrading the disk layout versions" on page 40.

31 Verify that all the file systems are working properly and data is intact.

See the cfsmount(1M) manual page.

32 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctrl status
```

33 If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctrl start
```

Upgrading the disk layout versions

SFCFS 5.0 MP2 supports disk layouts Versions 4, 5, 6 and 7 for locally mounted file systems and disk layouts Versions 6 and 7 for cluster mounted file systems. If you have cluster mounted file systems with disk layout versions lower than 6, then after upgrading to SFCFS 5.0 MP2, perform the following additional steps to prepare the file system for being mounted on all nodes of the cluster:

To upgrade the disk layout versions

1 Select one of the nodes of the cluster and mount the file system locally on this node. For example, mount it without the -o cluster option. Enter,

```
# mount -F vxfs block_device_path /mnt1
```

2 Current disk layout version on a file system can be found using

```
# fstyp -v char_device_path | grep version | \
    awk '{print $2}'
```

3 On the node selected in 1, incrementally upgrade the disk layout of this file system to layout Version 6 or layout Version 7. For example, if you had a cluster mounted file system of disk layout Version 4 while running with SFCFS 3.5 on HP-UX 11i Version 1, after upgrading to SFCFS 5.0 MP2, you would need to upgrade the disk layout to version 6 or version 7 incrementally as follows:

```
# vxupgrade -n 5 /mnt1
# vxupgrade -n 6 /mnt1
# vxupgrade -n 7 /mnt1
```

4 On the node selected in 1, after the disk layout has been successfully upgraded, unmount the file system.

```
# umount /mnt1
```

5 This file system can be mounted on all nodes of the cluster using `cfsmount`.

Upgrading Storage Foundation Cluster File System and High Availability software from a release 5.0 or later

Use this procedure to upgrade to 5.0 MP2 from 5.0 or later.

There are two ways to upgrade cluster nodes to the latest version of Storage Foundation Cluster File System: phased and full.

See "Phased upgrade for a Maintenance Pack" on page 42.

See "Full upgrade for a Maintenance Pack" on page 48.

An upgrade requires stopping cluster failover functionality during the entire procedure. The upgrade is performed in a number of stages depending on the type of upgrade you are performing.

You must have superuser (root) privileges to install the Veritas software.

You should also review the *Veritas Storage Foundation Release Notes* for important release information.

Caution: A phased upgrade procedure results in a system PANIC on configurations where LLT is configured over UDP. This issue is fixed in 5.0 MP1. This issue is specific to configurations where LLT is configured over UDP and not present in usual LLT Ethernet configurations. The full upgrade procedure should be used for upgrading from SFCFS 5.0 or SFCFS 5.0 RP1 to SFCFS 5.0 MP1 on configurations where LLT is configured over UDP.

Preparing to upgrade to the Maintenance Pack

If you are upgrading an installed Veritas Storage Foundation Cluster File System 5.0 version or from an earlier 5.0 Maintenance Pack, preserve the existing configuration information.

To preserve the existing configuration information, perform the following actions:

- Make a record of the mount points for VxFS file systems and VxVM volumes that are defined in the `/etc/fstab` file. You will need to recreate these entries in the `/etc/fstab` file on the freshly installed system.

- Before upgrading, ensure that you have made backups of all data that you want to preserve. In particular, you will need the information in files such as `/etc/fstab`. You should also run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands, and record the output from these. You may need this information to reconfigure your system after the upgrade.

- Use the `vxlicrep` command to make a record of the currently installed Veritas licenses.

Phased upgrade for a Maintenance Pack

A phased upgrade minimizes downtime by upgrading portions of the cluster, one at a time.

Although the entire cluster is offline for a shorter period than a full upgrade, this method requires command-line interaction and some manual configuration.

Each phase of the phased upgrade should be performed on more than one node of the cluster.

The stages of the phased upgrade procedure include the following steps:

- Freeze service group operations and stop cluster failover operations.

- Select a two or more nodes to upgrade, and leave a group of one or more nodes running.

- Take the selected group of nodes offline and prepare them for the upgrade.

- Upgrade the Veritas Storage Foundation Cluster File System and High Availability software on the selected group of nodes.

- Take the second group of nodes offline.

- Bring the first group of nodes online.

- Upgrade the second group of nodes.

- Bring the second group of nodes online and restart cluster failover services. The cluster is fully restored.

Performing the phased upgrade for a Maintenance Pack

This section describes how to perform a phased upgrade for a Maintenance Pack.

To freeze service group operations and stop cluster failover

1   Log in as superuser.

2   Verify that `/opt/VRTS/bin` is in your PATH so that you can execute all product commands.

3   From any node in the cluster, make the cluster configuration writable.

```
# haconf -makerw
```

4   Enter the following command to freeze high availability service group operations on each node:

```
# hasys -freeze -persistent node_name
```

5   Make the configuration read-only

```
# haconf -dump -makero
```

To select the nodes for the upgrade

1   Select one or more nodes to upgrade first.

2   Leave a group of one or more nodes running.

Upgrading the Veritas Storage Foundation Cluster File System and High Availability software to a Maintenance Pack

After the selected group of nodes is offline, the Veritas Storage Foundation Cluster File System software can be upgraded, using `installmp` for the Maintenance Pack.

To take the selected group of nodes offline and prepare them for the upgrade

1   Offline all SFCFS resources on nodes to be upgrade by running the following commands on one of the cluster nodes.

```
# hagrp -offline service_group -sys system01
# hagrp -offline service_group -sys system02
```

2   Stop cluster operations on each node in the group being upgraded, by entering the following command:

```
# hastop -local
```

3 On each node, use the following command to check if any VxFS file systems
or Storage Checkpoints are mounted:

```
# bdf -t vxfs
```

4 On each node in the cluster, unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name
# umount /filesystem
```

5 If you have created any Veritas Volume Replicator (VVR) replicated volume
groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a
  data volume contains a file system, unmount it.

- Use the vxrvg stop command to stop each RVG individually:

  ```
  # vxrvg -g diskgroup stop rvg_name
  ```

- On the Primary node, use the vxrlink status command to verify that all
  RLINKs are up-to-date:

  ```
  # vxrlink -g diskgroup status rlink_name
  ```

  To avoid data corruption, do not proceed until all RLINKs are up-to-date.

6 Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctrl status
```

7 If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctrl stop
```

8 If there are still disk groups that are imported at this time then proceed with
the remaining steps. Otherwise, skip to the procedure to upgrade the Veritas
software.

9 Stop activity to all VxVM volumes. For example, stop any applications such
as databases that access the volumes, and unmount any file systems that
have been created on the volumes.

10 On each node, stop all VxVM volumes by entering the following command
for each disk group:

```
# vxvol -g diskgroup stopall
```

To upgrade the Veritas Storage Foundation Cluster File System and High Availability software

1   Insert the appropriate media disc into your system's DVD-ROM drive.

2   Determine the block device file for the DVD drive:

    ```
    # ioscan -fnC disk
    ```

    Make a note of the device file as it applies to your system.

3   Create a directory in which to mount the software disc and mount the disc using the appropriate drive name. For example:

    ```
    # mkdir -p /mnt/cdrom
    # /usr/sbin/mount -F cdfs /dev/dsk/c3t2d0 /mnt/cdrom
    ```

4   Change to the top-level directory on the disc:

    ```
    # cd /mnt/cdrom
    ```

5   To upgrade the Storage Foundation Cluster File System, you must invoke the installmp command from one of your cluster nodes using the option that corresponds to your configuration:

    ■ To install on the local system only, enter the following command:

        ```
        # ./installmp
        ```

    ■ To install on more than one system using secure shell (ssh) utilities, enter the following command:

        ```
        # ./installmp node_name1 node_name2 ...
        ```

    ■ To install on more than one system using remote shell (rsh) utilities, enter the following command:

        ```
        # ./installmp node_name1 node_name2 ... -rsh
        ```

6   After the initial system checks are complete, press **Return** to start the requirement checks.

7   After the requirement checks are complete, press **Return** to start upgrading the packages. If you are upgrading multiple nodes, you have the option of upgrading them simultaneously. You will be prompted after the upgrade is complete.

8 When installation is complete, note the locations of the summary, log, and response files indicated by the installer.

9 Reboot the system (or systems).

Upgrading the remaining nodes

This section describes how to upgrade the remaining nodes.

Take the second group of nodes offline.

Bring the first group (with the newly installed patches) online.

Upgrade the second group of nodes.

To take the second group of nodes offline

1 Offline all SFCFS resources on nodes to be upgrade by running the following commands on one of the cluster nodes.

```
# hagrp -offline service_group -sys system01
# hagrp -offline service_group -sys system02
```

2 Stop cluster operations on each node in the second group being upgraded, by entering the following command:

```
# hastop -local
```

To bring the first group of nodes online

◆ Bring the first group of nodes online.

To upgrade the second group of nodes

1 To upgrade the second group of nodes, perform the upgrade of the Veritas Storage Foundation Cluster File System and High Availability software on the second group of nodes.

2 Then bring the second group of nodes online.

Bringing the upgraded group of nodes online

Use the following procedure to bring the upgraded group of nodes online.

To bring the upgraded group of nodes online

1   If you need to re-encapsulate and mirror the root disk on each of the nodes,
    follow the procedures in the "Administering Disks" chapter of the *Veritas
    Volume Manager Administrator's Guide*.

2   If necessary, reinstate any missing mount points in the `/etc/fstab` file on
    each node.

3   If any VCS configuration files need to be restored, stop the cluster, restore
    the files to the `/etc/VRTSvcs/conf/config` directory, and restart the cluster.

4   Make the VCS configuration writable again from any node in the upgraded
    group:

    ```
    # haconf -makerw
    ```

5   Enter the following command on each node in the upgraded group to unfreeze
    HA service group operations:

    ```
    # hasys -unfreeze -persistent node_name
    ```

6   Make the configuration read-only:

    ```
    # haconf -dump -makero
    ```

7   Bring the CVM service group online on each node in the upgraded group:

    ```
    # hagrp -online cvm -sys node_name
    ```

8   Restart all the volumes by entering the following command for each disk
    group:

    ```
    # vxvol -g diskgroup startall
    ```

9   If you have stopped any Veritas Volume Replicator (VVR) replicated volume
    groups (RVGs) on your system, restart each RVG:

    ```
    # vxrvg -g diskgroup start rvg_name
    ```

10 Enter the following command for each disk group:

    ```
    # vxdg -g diskgroup set activation=sw
    ```

11 Remount all VxFS file systems and Storage Checkpoints on all nodes:

```
# mount /filesystem
# mount /checkpoint_name
```

12 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctrl status
```

13 If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctrl start
```

Full upgrade for a Maintenance Pack

A full upgrade upgrades the product on the entire cluster and the cluster remains offline for the duration of the procedure. Minimal command-line interaction and some manual configuration are required.

The stages of the full upgrade procedure are:

■ Take all nodes in the cluster offline and install the software patches.

■ Bring all the nodes (with the newly installed patches) online to restart cluster failover services. The cluster is fully restored.

Performing the full upgrade to a Maintenance Pack

This section describes how to perform a full upgrade to a Maintenance Pack.

A full upgrade upgrades the product on the entire cluster and the cluster remains offline for the duration of the procedure. Minimal command-line interaction and some manual configuration are required.

To prepare for a full upgrade to a Maintenance Pack

1 Log in as superuser.

2 Verify that /opt/VRTS/bin is in your PATH so you can execute all product commands.

3 Stop high-availability cluster operations. This command can be executed from any node in the cluster, and stops cluster operations on all the nodes.

```
# hastop -all
```

4 Use the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# bdf -t vxfs
```

5 Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name
# umount /filesystem
```

6 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.

- Use the vxrvg stop command to stop each RVG individually:

  ```
  # vxrvg -g diskgroup stop rvg_name
  ```

- On the Primary node, use the vxrlink status command to verify that all RLINKs are up-to-date:

  ```
  # vxrlink -g diskgroup status rlink_name
  ```

  To avoid data corruption, do not proceed until all RLINKs are up-to-date.

7 Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctrl status
```

8 If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctrl stop
```

9 If there are still disk groups that are imported at this time then proceed with the remaining steps. Otherwise, skip to the procedure to upgrade the Veritas software.

10 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.

11 Stop all VxVM volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

12 To verify that no volumes remain open, use the following command:

```
# vxprint -Aht -e v_open
```

13 Continue to the procedure to upgrade the Veritas Storage Foundation Cluster File System and High Availability software.

To upgrade the Veritas Storage Foundation Cluster File System and High Availability software

1 Insert the appropriate media disc into your system's DVD-ROM drive.

2 Determine the block device file for the DVD drive:

```
# ioscan -fnC disk
```

Make a note of the device file as it applies to your system.

3 Create a directory in which to mount the software disc and mount the disc using the appropriate drive name. For example:

```
# mkdir -p /mnt/cdrom
# /usr/sbin/mount -F cdfs /dev/dsk/c3t2d0 /mnt/cdrom
```

4 Change to the top-level directory on the disc:

```
# cd /mnt/cdrom
```

5 To upgrade the Storage Foundation Cluster File System, you must invoke the installmp command from one of your cluster nodes using the option that corresponds to your configuration:

- To install on the local system only, enter the following command:

  ```
  # ./installmp
  ```

- To install on more than one system using secure shell (SSH) utilities, enter the following command:

  ```
  # ./installmp node_name1 node_name2 ...
  ```

- To install on more than one system using remote shell (RSH) utilities, enter the following command:

  ```
  # ./installmp node_name1 node_name2 ... -rsh
  ```

6 After the initial system checks are complete, press **Return** to start the requirement checks.

7 When installation is complete, note the locations of the summary, log, and response files indicated by the installer.

8 Shut down and reboot the systems.

To bring the upgraded cluster online and restore components

1 If you need to re-encapsulate and mirror the root disk on each of the nodes, follow the procedures in the "Administering Disks" chapter of the *Veritas Volume Manager Administrator's Guide*.

2 If necessary, reinstate any missing mount points in the `/etc/fstab` file on each node.

3 If any VCS configuration files need to be restored, stop the cluster, restore the files to the `/etc/VRTSvcs/conf/config` directory, and restart the cluster.

4 Restart all the volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup startall
```

5 If you have stopped any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, restart each RVG:

```
# vxrvg -g diskgroup start rvg_name
```

6 Enter the following command for each disk group:

```
# vxdg -g diskgroup set activation=sw
```

7 Remount all VxFS file systems and Storage Checkpoints on all nodes:

```
# mount /filesystem
# mount /checkpoint_name
```

8 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctrl status
```

9 If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctrl start
```

Verifying the configuration files for Storage Foundation Cluster File System

You can inspect the contents of the configuration files that were installed and modified after a successful installation process. These files reflect the configuration based on the information you supplied.

To verify the configuration files

1   Log in as superuser to any system in the cluster.

2   Set up your environment PATH variable.

    # **export PATH=$PATH:/sbin:/usr/sbin:/opt/VRTS/bin**

Low Latency Transport configuration files

The following files are required by the VCS communication services for Low Latency Transport (LLT).

/etc/llthosts

The file llthosts(4M) is a database, containing one entry per system, that links the LLT system ID (in the first column) with the LLT host name. This file is identical on each system in the cluster.

For example, the file /etc/llthosts contains entries that resemble:

```
0     system01
1     system02
```

/etc/llttab

The file llttab(4M) contains information that is derived during installation and used by the utility lltconfig(1M). After installation, this file lists the network links that correspond to the specific system.

For example, the file /etc/llttab contains entries that resemble:

```
set-node system01
set-cluster 100

link lan1 lan:1 - ether - -
link lan2 lan:2 - ether - -
```

The first line identifies the local system name. The second line identifies the cluster (that is, the cluster ID you entered during installation). The next two lines, beginning with the `link` command, identify the two network cards used by the LLT protocol.

See the `llttab`(4M) manual page.

Checking Low Latency Transport operation

Use the `lltstat` command to verify that links are active for LLT. This command returns information about the links for LLT for the system on which it is typed. See the `lltstat`(1M) manual page.

In the following example, `lltstat -n` is typed on each system in the cluster.

To check LLT operation

1 Log into system01.

   # **lltstat -n**

   Output resembles:

   ```
       LLT   node   information:
        Node                           State          Links
   *   0    system01               OPEN               2
       1    system02               OPEN               2
   ```

2 Log into system02.

    # **lltstat -n**

   Output resembles:

   ```
       LLT   node   information:
        Node                           State          Links
       0    system01               OPEN               2
   *   1    system02               OPEN               2
   ```

   Each system has two links and that each system is in the OPEN state. An asterisk (*) denotes the system on which the command is typed.

   With LLT configured correctly, the output of `lltstat -n` shows all of the systems in the cluster and two links for each system. If the output shows otherwise, you can use the verbose option of `lltstat`. For example, type `lltstat -nvv | more` on a system to view additional information about LLT. In the following example, `lltstat -nvv | more` is typed on a system in a two-node cluster.

3   Log into system01.

```
# lltstat -nvv | more
```

Output resembles:

```
Node          State        Link         Status        Address

*0 system01 OPEN           lan1         UP            08:00:20:93:0E:34
                           lan2         UP            08:00:20:93:0E:34
 1 system02  OPEN          lan1         UP            08:00:20:8F:D1:F2
                           lan2         DOWN          08:00:20:8F:D1:F2

 2           CONNWAIT
                           lan1         DOWN
                           lan2         DOWN
 .
 .
 .
 31          CONNWAIT
                           lan1         DOWN
                           lan2         DOWN
```

The output lists 32 nodes. It reports on the two cluster nodes, system01 and
system02, plus non-existent nodes. For each correctly configured system,
the information shows a state of OPEN, a status for each link of UP, and an
address for each link. However, in the example above, the output shows that
for node system02, the private network may have failed earlier, or the
information in /etc/llttab may be incorrect.

To obtain information about the ports open for LLT, type lltstat -p on any
system. In the following example, lltstat -p is typed on one system in the
cluster.

4 Log into system01.

lltstat -p

Output resembles:

```
LLT port information:
 Port    Usage    Cookie
 0       gab      0x0
         opens:          0 1 3 4 5 6 7 8 9 10 11 12 13...
         connects:       0 1
```

The two systems with node ID's 0 and 1 are connected.

See

Group Membership and Atomic Broadcast configuration files

The following files are required by the VCS communication services for Group
Membership and Atomic Broadcast (GAB).

/etc/gabtab

After installation, the file /etc/gabtab contains a gabconfig(1M) command that
configures the GAB driver for use.

The file /etc/gabtab contains a line that resembles:

```
/sbin/gabconfig -c -nN
```

where the -c option configures the driver for use and -n*N* specifies that the cluster
will not be formed until at least *N* systems are ready to form the cluster. *N* is the
number of systems in the cluster.

Checking Group Membership and Atomic Broadcast operation

This section describes how to check GAB operation.

To check GAB operation

◆ Enter the following command on each node in the cluster.

```
# gabconfig -a
```

If GAB is operational, the following output displays with GAB port membership
information:

```
GAB Port Memberships
==================================================================
Port a gen 1bbf01 membership 01
Port b gen 1bbf06 membership 01
Port f gen 1bbf0f membership 01
Port h gen 1bbf03 membership 01
Port v gen 1bbf0b membership 01
Port w gen 1bbf0d membership 01
```

If GAB is not operational, the following output display with no GAB port
membership information:

```
GAB Port Memberships
==================================================================
```

See the *Veritas Cluster Server User's Guide*.

Checking cluster operation

This section describes how to check cluster operation.

To check cluster operation

1 Enter the following command on any system:

```
# hastatus -summary
```

The output for an SFCFS HA installation resembles:

```
-- SYSTEM STATE
-- System                    State               Frozen

A  system01                 RUNNING             0
A  system02                 RUNNING             0

-- GROUP STATE
-- Group      System      Probed AutoDisabled  State

B  cvm        system01    Y      N             ONLINE
B  cvm        system02    Y      N             OFFLINE
```

If the State value is running, VCS is successfully installed and running on that node. The group state lists the cvm group, which is online on system01 and offline on system02.

See the hastatus(1M) manual page.

See the *Veritas Cluster Server User's Guide*.

2 Enter the following command on any systems:

```
# hasys -display
```

See the *Veritas Cluster Server User's Guide*.

For more information on the hasys -display command, see the hasys(1M) manual page.

The example shows the output of system01. The list continues with similar information for system02 (not shown) and any other systems in the cluster. The output should be similar on each system.

Table 1-4 shows on each system, the output should be similar:

Table 1-4   System sample output for hasys -display

#System	Attribute	Value
system01	AgentsStopped	0
system01	AvailableCapacity	1

Table 1-4   System sample output for hasys -display *(continued)*

#System	Attribute	Value
system01	Capacity	1
system01	ConfigBlockCount	54
system01	ConfigCheckSum	29776
system01	ConfigDiskState	CURRENT
system01	ConfigFile	/etc/VRTSvcs/conf/config
system01	ConfigInfoCnt	0
system01	ConfigModDate	Tues June 25 23:00:00 2006
system01	CurrentLimits	
system01	DiskHbStatus	
system01	DynamicLoad	0
system01	Frozen	0
system01	GUIIPAddr	
system01	LLTNodeId	0
system01	Limits	
system01	LoadTimeCounter	1890
system01	LoadTimeThreshold	600
system01	LoadWarningLevel	80
system01	MajorVersion	2
system01	MinorVersion	0
system01	NodeId	0
system01	OnGrpCnt	1
system01	ShutdownTimeout	60
system01	SourceFile	./main.cf
system01	SysName	system01

Table 1-4   System sample output for hasys -display *(continued)*

#System	Attribute	Value
system01	SysState	RUNNING
system01	SystemLocation	
system01	SystemOwner	
system01	TFrozen	0
system01	TRSE	0
system01	UpDownState	Up
system01	UserInt	0
system01	UserStr	

Synchronizing time on Cluster File Systems

SFCFS requires that the system clocks on all nodes are synchronized using some external component such as the Network Time Protocol (NTP) daemon. If the nodes are not in sync, timestamps for change (ctime) and modification (mtime) may not be consistent with the sequence in which operations actually happened.

Configuring VCS for Storage Foundation Cluster File System

Configuring VCS means conveying to the VCS engine the definitions of the cluster, service groups, resources, and resource dependencies. VCS uses two configuration files in a default configuration:

■ The main.cf file defines the entire cluster.

■ The types.cf file defines the resource types.

By default, both files reside in the directory /etc/VRTSvcs/conf/config. Additional files similar to types.cf may be present if agents have been added, such as OracleTypes.cf.

In a VCS cluster, the first system to be brought online reads the configuration file and creates an internal (in-memory) representation of the configuration. Systems brought online after the first system derive their information from systems running in the cluster. You must stop the cluster while you are modifying the files

from the command line. Changes made by editing the configuration files take effect when the cluster is restarted. The node on which the changes were made should be the first node to be brought back online.

main.cf file

The VCS configuration file `main.cf` is created during the installation procedure. After installation, the `main.cf` file contains the base definitions of the cluster and its nodes. Additionally, the file `types.cf` listed in the include statement defines the bundled agents for VCS resources.

See the *Veritas Cluster Server User's Guide*.

A typical VCS configuration file for SFCFS file resembles:

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"

cluster sfcfs_1 (
        HacliUserLevel = COMMANDROOT
        )

system thor150 (
        )

system thor151 (
        )

group cvm (
        SystemList = { thor150 = 0, thor151 = 1 }
        AutoFailOver = 0
        Parallel = 1
        AutoStartList = { thor150, thor151 }
        )

        CVMCluster cvm_clus (
                CVMClustName = sfcfs_1
                CVMNodeId = { thor150 = 0, thor151 = 1 }
                CVMTransport = gab
                CVMTimeout = 200
                )

        CVMVxconfigd cvm_vxconfigd (
```

```
                  Critical = 0
                  CVMVxconfigdArgs = { syslog }
                  )

          cvm_clus requires cvm_vxconfigd



          // resource dependency tree
          //
          //      group cvm
          //      {
          //      CVMCluster cvm_clus
          //          {
          //          CVMVxconfigd cvm_vxconfigd
          //          }
          //      }
```

Storage Foundation Cluster File System HA Only

If you configured VCS Cluster Manager (Web Console), a service group, "ClusterService," was created that includes IP, Process, and Notifier resources. These resources were configured according to information you provided during the installation procedure. A resource dependency was also created.

VCS application failover services

If you installed SFCFS HA, you can begin implementing the application monitoring failover services provided by the Veritas Cluster Server. Information about setting up VCS services is beyond the scope of this document.

See the *Veritas Cluster Server* documentation.

Uninstalling Storage Foundation Cluster File System

If you need to uninstall SFCFS software. Use the `uninstallsfcfs` script.

To uninstall SFCFS HA

1   Log in as superuser.

2   Stop the cluster:

    # hastop -all

    Do not use the `hastop -force` command to stop VCS.

3 Change directory to `/opt/VRTS/install`:

`# `**`cd /opt/VRTS/install`**

4 Run the `uninstallsfcfs` command to uninstall SFCFS. The `uninstallsfcfs` script uses ssh to communicate with remote nodes as default:

`# `**`./uninstallsfcfs`**

If you want to use rsh you must specify on the command line:

`# `**`./uninstallsfcfs -rsh`**

5 Enter the system names to uninstall SFCFS.

```
Enter the system names separated by spaces on which to
uninstall SFCFS: system01 system02
```

6 Enter **y** to uninstall SFCFS.

```
Are you sure you want to uninstall SFCFS? [y,n,q] (y)
```

7 Reboot the systems on which SFCFS is uninstalled after successful uninstallation.

Uninstalling the 5.0 MP2 patches

The following procedure removes the patches if you want to uninstall 5.0 MP2.

To remove the 5.0 MP2 patches

1 Log in as superuser.

2 Stop the cluster:

`# `**`hastop -all`**

3 On all the nodes, stop VxFEN:

`# `**`/sbin/init.d/vxfen stop`**

4 On all the nodes, stop the currently running VxPAL agents (see `vxpalctrl` (1M)):

■ Stop the storage agent:

`# `**`/opt/VRTSobc/pal33/bin/vxpalctrl -a StorageAgent -c stop`**

■ Check the status of the storage agent:

/opt/VRTSobc/pal33/bin/vxpalctrl -a StorageAgent -c status

■ Stop the action agent:

/opt/VRTSobc/pal33/bin/vxpalctrl -a actionagent -c stop

■ Check the status of the action agent:

/opt/VRTSobc/pal33/bin/vxpalctrl -a actionagent -c status

■ Stop the gridnode agent:

/opt/VRTSobc/pal33/bin/vxpalctrl -a gridnode -c stop

■ Check the status of the gridnode agent:

/opt/VRTSobc/pal33/bin/vxpalctrl -a gridnode -c status

5 On all the nodes, remove all the patches using the swremove command:

swremove -x autoreboot=true *patch_name1*, *patch_name2* ...

See "Storage Foundation Cluster File System patches" on page 63.

Symantec recommends that all the patches installed during 5.0 MP2 installation
be removed through a single command line. The system automatically reboots
after removing the patches.

Storage Foundation Cluster File System patches

The following are the patches that are installed as a part of the SFCFS 5.0 MP2
installation:

Table 1-5   Storage Foundation Cluster File System patches

Patch ID	Patch Name
PHCO_38997	VRTSobc33 Command Patch
PHCO_38381	VRTSob Command Patch
PHCO_38383	VRTSobgui Command Patch
PHCO_38385	VRTSccg Command Patch

Table 1-5 Storage Foundation Cluster File System patches *(continued)*

Patch ID	Patch Name
PHCO_35217	VRTSmh Command Patch
PHCO_38384	VRTSaa Command Patch
PVCO_03698	VRTSsmf Command Patch
PVCO_03697	SYMClma Command Patch
PHNE_38739	VRTSllt Command Patch
PHNE_38738	VRTSgab Command Patch
PHCO_38740	VRTSvxfen Command Patch
PHKL_38743	VRTSvxfen Kernel Patch
PVCO_03797	VRTSvcs/VRTSvcsag Command Patch
PHSS_35963	VRTSjre15 Command Patch
PVCO_03678	VRTScscm Command Patch
PVCO_03690	VRTScscw Command Patch
PVCO_03680	VRTScssim Command Patch
PVCO_03689	VRTScmcs Command Patch
PVCO_03686	VRTScmccc Command Patch
PHCO_38830	VRTSvxvm Command Patch
PHKL_38829	VRTSvxvm Kernel Patch
PHCO_38859	VRTSfspro Command Patch
PHCO_38831	VRTSvmpro Command Patch
PHCO_37077	VRTSdcli Command Patch
PHCO_38981	VRTSalloc Command Patch
PHCO_38834	VRTSddlpr Command Patch
PVCO_03671	VRTSvrpro Command Patch
PVCO_03672	VRTSvrw GUI Patch
PHCO_38836	VRTSfsman Command Patch

Table 1-5   Storage Foundation Cluster File System patches *(continued)*

Patch ID	Patch Name
PHCO_38850	VRTSvxfs Command Patch
PHKL_38794	VRTSvxfs Kernel Patch
PHKL_38796	VRTSglm Kernel Patch
PVCO_03850	VRTScavf Command Patch

Software limitations

Veritas Storage Foundation Cluster File System software limitations in the 5.0 release are listed in the *Veritas Storage Foundation Cluster File System 5.0 Release Notes.*

The *Veritas Storage Foundation Cluster File System 5.0 Release Notes* can be viewed at the following URL:

http://entsupport.symantec.com/docs/283716

The Veritas Storage Foundation, Veritas Volume Manager, Veritas File System, Veritas Volume Replicator, and Veritas Storage Foundation for Databases software limitations are listed in the Veritas Storage Foundation 5.0 MP2 Release Notes.

See the *Veritas Storage Foundation 5.0 MP2 Release Notes.*

5.0 MP2 software limitations

There are no new Veritas Storage Foundation Cluster File System software limitations in this 5.0 MP2 release.

Fixed issues

Veritas Storage Foundation Cluster File System fixed issues in the 5.0 release are listed in the *Veritas Storage Foundation Cluster File System 5.0 Release Notes.*

The *Veritas Storage Foundation Cluster File System 5.0 Release Notes* can be viewed at the following URL:

http://entsupport.symantec.com/docs/283716

The Veritas Storage Foundation, Veritas Volume Manager, Veritas File System, Veritas Volume Replicator, and Veritas Storage Foundation for Databases fixed issues are listed in the *Veritas Storage Foundation 5.0 MP2 Release Notes.*

See the *Veritas Storage Foundation 5.0 MP2 Release Notes.*

5.0 MP2 fixed issues

The following are Veritas Storage Foundation Cluster File System fixed issues in 5.0 MP2 release:

Table 1-6   5.0 MP2 Storage Foundation Cluster File System fixed issues

Incident	Description
1510791	After an open("..") failed with ENOENT because the parent directory was removed, subsequent open("..") still failed with ENOENT even if the parent directory was re-created. A negative DNLC entry is now not added if the directory inode is marked for removal.
1507758	An oversized transaction processing goes into an infinite loop and leads to a system hang with a vx_trancommit() function failure. Reservation sizes are now increased only gradually after vx_trancommit() failures.
1470927	While deleting a large file (~>2GB) on a cluster mounted filesystem having a checkpoint, the system stops responding. The issue occurs because large file data buffers unnecessarily get invalidated. Only the meta-data buffers now get invalidated.
1465699	A Data page fault panic occurred in the vx_statvfs() function. The vx_statvfs code is now modified to lock and validate the VFS structures.
1465693	Systems experienced a map lock contention where the CPU was unnecessarily relinquished in cases where the logs were flushed asynchronously. The CPU is now not relinquished in cases where the logs are flushed asynchronously.

Table 1-6   5.0 MP2 Storage Foundation Cluster File System fixed issues
(continued)

Incident	Description
1425413, 1425405	The `qiomkfile` file had incorrect permissions set. Permissions on the `qiomkfile` are now corrected.
1417404	The `vxupgrade` command allowed `/` and `/stand` filesystems to be upgraded to disk layout Version greater than 5. A check is now added to prevent the `vxupgrade` command from upgrading `/` and `/stand` filesystems to disk layout Version greater than 5.
1274326	A livelock condition occurred while updating an inode. This caused the Cluster File System (CFS) to hang. An `ILOCK` is now dropped and then reacquired. The `ILOCK` prevents the livelock condition.
1274317	The intent log replay during a full `fsck` operation took approximately 2 to 3 hours to execute. Tail pointers are now added to the linked lists. As a result, the `fsck` utility takes less time to execute.
1274315	Internal stress testing on a cluster filesystem causes an assertion failure while processing inactive inodes. Inactive processing does not now occur more than once on an attribute inode. A new pointer keeps a track of the inactive processing to prevent reprocessing the attribute inodes.

Table 1-6   5.0 MP2 Storage Foundation Cluster File System fixed issues
(continued)

Incident	Description
1040634	While upgrading disk layout from Version 5 to Version 6 using the `vxupgrade` command, the file system would freeze with a file table overflow. The background threads, which executed the processing of inactive inodes for the upgraded filesystem are now removed. Instead, the existing upgrade threads perform the inactive inodes processing and prevent the inodes from building up in the filesystem memory. The unbound attribute-ilist reorg inodes no longer build up in the filesystem memory and are instead re-used. The vnodeops is no longer checked. Therefore, inodes can now be reused.
1029522	`sar -v` reported an incorrect `inod-sz` and other tools that called `pstat()` with the `PSTAT_DYNAMIC` option returned an incorrect inode cache size. The `ioctl()` value of `VX_GET_NINODE` in the 4.1 release of VxFS is now changed to match the value in release 3.5.
896446	LM noise.fullfsck hit assert f:vx_msgprint:ndebug via vx_disable for build 5.0A53 on OPTERON + SLES9 SP3 RC4-vsmp(64Bit).

5.0 MP1 fixed issues

The following are Veritas Storage Foundation Cluster File System fixed issues in 5.0 MP1 release:

Table 1-7   5.0 MP1 Storage Foundation Cluster File System fixed issues

Incident	Description
156186	In environments using LDAP or NIS+, the CFSfsckd agent scripts could hang or time out because of unresponsive network-based name services occurring due to incorrect options included in the ps command. The problem has been resolved by setting the correct options for the ps command to avoid name lookups over the network.
647359	Under certain conditions, the cmhaltcl (1M) command was failing to halt the cluster nodes in a ServiceGuard cluster when invoked with the -f option. This was occurring due to the incorrect freeing of CFS response messages during the retry of a previously sent message in case of an error. The problem has been resolved by freeing the cluster node's response message in case of an error so that the failed message does not get allocated again.
704284	The fsclustadm (1M) was continuously returning a device busy message when using the cfsdeinit command line option, in a ServiceGuard cluster. During a forced unmount of a CFS file system the hold count was not being correctly updated causing an EBUSY error status to be returned to fsclustadm (1M). The problem has been resolved by correctly updating the hold count during inode inactivation.
771108	Oracle 10gr1 database creation on raw volumes with ODM was failing. The problem has been resolved by adding a new permission bit to the ODM interface.
797453	Performance enhancements were required in the handling of inherited extended attributes. The problem has been resolved by ensuring that LCT buffers are written to disk only when necessary instead of every time they are modified.

Table 1-7   5.0 MP1 Storage Foundation Cluster File System fixed issues
(continued)

Incident	Description
799189	The Oracle instance was crashing because ODM was returning `ENOMEM`. The problem has been resolved by changing the arena allocation call to use the `WAITOK` flag.
802327	A Data Page Fault was being encountered in a fast lookup path for NFS access when a node joins the cluster. The problem has been resolved by ensuring that the code paths that provide fast access for NFS requests, initialize the vfs pointer whenever it has not yet been initialized.
816497	The direct access of a 64-bit data item in a CFS message was causing an alignment trap. After a cluster node was rebooted, a remaining cluster node was crashing in `vx_do_rbdele_resp()` due to the alignment trap. The problem has been resolved by avoiding direct access of the 64-bit data item in the CFS message by using `bzero()` and `bcopy()` to move data in and out of the CFS message buffer.
817879	A blocking file lock request over NFS was resulting in a hung cluster, for a file that is residing on a cluster mounted file system. The problem has been resolved by ensuring that the file locking code in CFS handles the blocking file lock requests over NFS appropriately.
847819	A Data Page Fault was being encountered in `vx_find_fsparams()` due to a missing check for an unmounted file system. The problem has been resolved by ensuring that the file system has not been unmounted before accessing the in-core structures of the file system in `vx_async_getdele()`.

Known issues

Veritas Storage Foundation Cluster File System known issues in the 5.0 release are listed in the *Veritas Storage Foundation Cluster File System 5.0 Release Notes.*

The *Veritas Storage Foundation Cluster File System 5.0 Release Notes* can be viewed at the following URL:

http://entsupport.symantec.com/docs/283716

The Veritas Storage Foundation, Veritas Volume Manager, Veritas File System, Veritas Volume Replicator, and Veritas Storage Foundation for Databases known issues are listed in the *Veritas Storage Foundation 5.0 MP2 Release Notes.*

See the *Veritas Storage Foundation 5.0 MP2 Release Notes.*

5.0 MP2 known issues

The following are Veritas Storage Foundation Cluster File System known issues in 5.0 MP2 release:

Warning messages may appear in the swinstall.log and swremove.log files (1529476)

Warning messages may appear in the `swinstall.log` files after you have installed the product using the common product `installer` script. For example:

```
WARNING: The dependencies for fileset
PHCO_38836.VXFS-ENG-A-MAN,r=1.0"cannot be resolved (see
previous lines).
The operation on this fileset will still be attempted
even though there are unresolved dependencies because the
"enforce_dependencies" option is set to "false".

WARNING: New Install PHCO_38836.VXFS-ENG-A-MAN,r=1.0
WARNING: 1 of 1 filesets had Warnings.
WARNING: The Analysis Phase had warnings. See the above
output for details.

WARNING: The dependencies for fileset
"PHCO_38850.VXFS-RUN,r=1.0"cannot be resolved (see
previous lines). The operation on this fileset will still
be attempted even though there are unresolved
dependencies because the "enforce_dependencies" option is
set to "false".
```

```
* Summary of Analysis Phase:
WARNING: New Install PHCO_38850.VXFS-RUN,r=1.0
WARNING: 1 of 3 filesets had Warnings.
* 2 of 3 filesets had no Errors or Warnings.
WARNING: The Analysis Phase had warnings. See the above
output for details.

WARNING: The dependencies for fileset
"PHKL_38829.VXVM-KRN,r=1.0" cannot be resolved (see
previous lines). The operation on this fileset will still
be attempted even though there are unresolved
dependencies because the "enforce_dependencies" option is
set to "false".

* Summary of Analysis Phase:
WARNING: New Install PHKL_38829.VXVM-KRN,r=1.0
WARNING: 1 of 1 filesets had Warnings.
WARNING: The Analysis Phase had warnings. See the above
output for details.

WARNING: The dependencies for fileset
"PHCO_38830.VXVM-RUN,r=1.0" cannot be resolved (see
previous lines). The operation on this fileset will still
be attempted even though there are unresolved
dependencies because the "enforce_dependencies" option is
set to "false".

* Summary of Analysis Phase:
WARNING: New Install PHCO_38830.VXVM-RUN,r=1.0
WARNING: 1 of 4 filesets had Warnings.
* 3 of 4 filesets had no Errors or Warnings.
WARNING: The Analysis Phase had warnings. See the above
output for details.
```

Warning messages may appear in the `swremove.log` file while removing the patches. For example:

```
* The fileset
"PHCO_38831.VMPRO-PRG,l=/opt/VRTSvmpro,r=1.0" requires
the selected fileset "PHKL_38829.VXVM-KRN,l=/,r=1.0"
as a corequisite.

WARNING: The dependencies for fileset
"PHKL_38829.VXVM-KRN,l=/,r=1.0"cannot be resolved (see
```

```
previous lines). The operation on this fileset will still
be attempted even though there are unresolved
dependencies because the "enforce_dependencies" option is
set to "false".
* The fileset
"PHCO_38831.VMPRO-PRG,l=/opt/VRTSvmpro,r=1.0" Storage
Foundation Cluster File System Release Notes 71
Known issues
requires the selected fileset
"PHCO_38830.VXVM-RUN,l=/,r=1.0"
as a corequisite.

WARNING: The dependencies for fileset
"PHCO_38830.VXVM-RUN,l=/,r=1.0"
cannot be resolved (see previous lines).

The operation on this fileset will still be attempted
even
though there are unresolved dependencies because the
"enforce_dependencies" option is set to "false".
* Summary of Analysis Phase:

WARNING: Remove PHKL_38829.VXVM-KRN,l=/,r=1.0
WARNING: Remove PHCO_38830.VXVM-RUN,l=/,r=1.0
WARNING: 2 of 46 filesets had Warnings.
* 44 of 46 filesets had no Errors or Warnings.
WARNING: The Analysis Phase had warnings. See the above
output for details.
```

These warning messages are harmless and may be ignored. The required patches
are correctly installed despite these warning messages.

During the install of SFCFS from DVD media, the installer may display an install failed error messages for VRTSvxvm package (1527735)

When you run the install of SFCFS and SF from DVD media, the installer may
display an install failed error messages for VRTSvxvm package. The installer may
be intervened by a timeout (600 seconds) during the VRTSvxvm package installation
from a DVD media depending on the speed of DVD driver.

Workaround

Use `installer -timeout 1800`. The timeout default value (600 seconds) will be overridden to 1800 seconds.

Information about a managed host is not updated in the Central Server GUI (1519528)

After upgrade from SFCFS 4.1MP2 to 5.0MP2, information about a managed host is not updated in the Central Server GUI.

Workaround

The following steps can be used to push the latest patch on the managed host after which the host information should be discovered correctly.

1   Remove the `UNOF` patch of `VRTSmh` on the managed host. For example:

    # **swremove UNOF_36026**

2   From the central server, push the latest available patch to the managed host using the `vxdompackage` command. For example:

    # **/opt/VRTScs/adm/vxdompackage add --hostfile **
    **/hostfile --imagefile **
    /opt/VRTScs/adm/vxdom_images/hpux/VRTSMH_1.1.tar

where *hostfile* is the file containing the managed host name.

Task Assistant tab is not available in the VEA GUI (1520408)

After upgrading from SFCFS 4.1 to SFCFS 5.0 the Task Assistant tab is not available in VEA GUI. Use the System tab to perform operations using VEA GUI.

vxinstalltemplate binary of VRTSalloc core dumps (1522127)

The vxinstalltemplate binary of VRTSalloc core dumps during 5.0MP2 SFCFS stack installation. A core file with name `core.vxinstalltemplate` will be generated.

Workaround

To resolve this issue, run the following command:

    # **/usr/sbin/vxtemplate -r -C -d **
    /opt/VRTSalloc/config/alloc_capabilities.txt install

5.0 MP1 known issues

The following are Veritas Storage Foundation Cluster File System known issues in 5.0 MP1 release:

Oracle-HP (840486)

Problems uninstalling or upgrading Veritas Storage Foundation for Oracle when Veritas Storage Foundation Cluster File System is installed on the same system.

If Veritas Storage Foundation for Oracle and Veritas Storage Foundation Cluster File System are installed on the same machine, do not use the installer to uninstall if you are planning to uninstall only one product.

If you want to uninstall the product, you must uninstall the Veritas Storage Foundation for Oracle packages manually.

To uninstall the Veritas Storage Foundation for Oracle packages

1   Review the uninstallation requirements in the *Veritas Storage Foundation Installation Guide.*

2   Stop the repository database and unmount the repository volume.

In a stand-alone configuration:	Stop the database repository:
	# **/opt/VRTSdbcom/bin/sfua_db_config -o stopdb**
	Unmount the database repository:
	# **/opt/VRTSdbcom/config/sfua_rep_mount stop**
In an HA configuration:	Stop VCS processes on either the local system or all systems.
	To stop VCS processes on the local system:
	# **hastop -local**
	To stop VCS processes on all systems:
	# **hastop -all**

3   Remove the Veritas Storage Foundation for Oracle packages using the swremove command.

    # **swremove VRTSorgui VRTSdbed VRTSdbcom VRTSdbdoc**

If Veritas Storage Foundation for Oracle and Veritas Storage Foundation Cluster File System are installed on the same machine and you are upgrading both products, use the installer to upgrade Veritas Storage Foundation Cluster File System first. Then, use the installer to upgrade Veritas Storage Foundation for Oracle.

If the second upgrade fails, remove the Veritas Storage Foundation for Oracle packages as described above, then run the installer to upgrade Veritas Storage Foundation for Oracle.